

Signalling and obfuscation for congestion control

Jakub Mareček, Robert Shorten, Jia Yuan Yu*
 IBM Research – Ireland
 Technology Campus Damastown, Dublin 15, Ireland

May 6, 2016

Abstract

We aim to reduce the social cost of congestion in many smart city applications. In our model of congestion, agents interact over limited resources after receiving signals from a central agent that observes the state of congestion in real time. Under natural models of agent populations, we develop new signalling schemes and show that by introducing a non-trivial amount of uncertainty in the signals, we reduce the social cost of congestion, i.e., improve social welfare. The signalling schemes are efficient in terms of both communication and computation, and are consistent with past observations of the congestion. Moreover, the resulting population dynamics converge under reasonable assumptions.

1 Introduction

The study of “Smarter Cities” and “Smarter Planet” provides a host of new and challenging control engineering problems. Many of these problems can be cast in a congestion control framework, where a large number of agents such as people, cars, or consumers compete for a limited resource. Examples of such problems include consumers competing for electricity supply; road users competing for space in the roads; users of bike (cars) sharing systems competing for access to bikes (cars); and pedestrians competing for access. Within each such problem, there are many variants. Pedestrians could be, for example, workers arriving at work through a common gate, people leaving a stadium, waiting to check in at an airport, or in an emergency. Each variant may have additional constraints, but certain key features remain the same.

Addressing these problems is hugely challenging, even without considering the issues of scale and limited communication. Resource utilisation should be minimised while delivering a certain quality of service to individual agents. Congestion is often caused by bursty arrivals, rather than the capacity and quality-of-service constraints *per se*. Prediction systems to alleviate congestion (informing customers of parking spaces, for example) create complicated

*Email: { jakub.marecek, robshort, jiayuan }@ie.ibm.com

feedback systems that are difficult to model and control. This latter issue of prediction and optimisation under feedback clearly opens a wide area of research, with links to reinforcement learning, adversarial game theory, and closed-loop identification at scale.

An example of congestion due to synchronised demand is the well-known flapping effect in road networks, parking lots, and bike-sharing stations. When presented with alternative choices of resources, agents who choose greedily based on past observations cause a congestion to oscillate between the resources over time. The flapping effect occurs due to actuation delays and the feedback effect between agents’ observations and actions: the agents’ choices are based on *past* observations, but affect the *future* state of congestion. Our principal motivation is to develop methods to break up this effect by taking into account these feedback issues and actuation delays. The starting point for our work is a simple model of agent-induced congestion.

Our objective in this paper is to take a first step toward addressing some of these problems. We wish to develop easy-to-implement algorithms, i.e. not requiring much inter-agent communication or dedicated infrastructure, with the objective of “de-synchronising” agents’ actions by providing them with signals, whereby achieving a temporal load balancing over networks. Load-balancing ideas in this direction have been recently suggested in a variety of applications [17, 15, 16]. However, the work presented here goes far beyond what has been proposed in several ways. First, it suggests non-trivial signalling schemes, as opposed to the simple randomisation and differential pricing ideas. Second, it takes into account heterogeneous agent behaviour and actions. Furthermore, our algorithms are provably scalable and can be analysed in a simple stochastic setting, whereby yielding provable behaviour even in the case of large communication and actuation delays.

Specifically, we model a congestion problem as a multiagent system evolving over time, where the agents follow natural policies. Roughly speaking, we show that by varying the amount of uncertainty agents face, one can reduce congestion in a controlled manner, and that a desirable state is arrived at asymptotically. Our investigation proceeds as follows. First, we fix the congestion cost functions, as well as the policies of each agent in a heterogeneous population. Then, we vary the parameters controlling the distribution of the random signals sent to the agents and evaluate the resulting social cost. We repeat this investigation for two signalling schemes, and show that the social cost is optimised by introducing uncertainty. Although our results are not difficult to derive, they are – to the best of our knowledge – both novel and useful in a smart city context.

The paper is organised as follows. After describing the model of congestion in Section 2, we set our work in context with the existing works in Section 3. Next, we present in two sections our main results for two models of signalling and agent-response. Section 4 considers a scalar signalling scheme and agents with different actuation delays. Section 5 considers a broadcast interval signalling scheme and agents with different risk aversion. In each of these two sections, we show through theoretical and empirical analysis that withholding a non-trivial amount of information from the agents reduce the overall social cost. Finally,

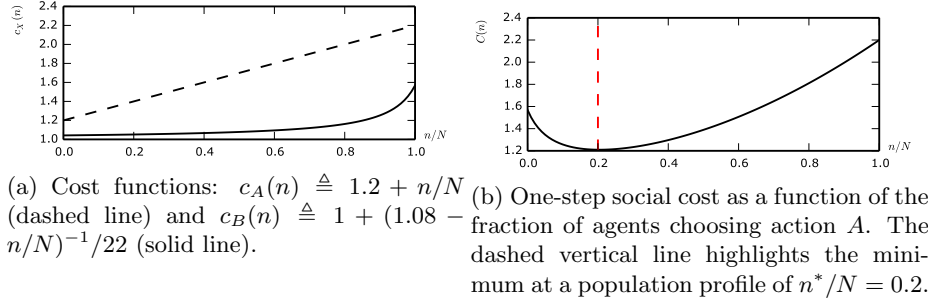


Figure 1: The setting of the running example. Despite the fact that the curves do not intersect, the optimum of the social cost is achieved when both actions are used – with a 0.2-fraction of agents choosing action A .

we present open questions in Section 6.

2 Model

We consider a model of congestion, where a population of N agents $\{1, \dots, N\}$ is confronted with two alternative choices at discrete time steps. Note this situation is widespread and is generalised to the case if multiple choices (more than two) in a straightforward manner¹. Our approach is to model congestion using probabilistic techniques. In this context alternative agent actions are denoted by $\{A, B\}$, and time steps are denoted by $t = 1, 2, \dots$. Then, the random variable a_t^i denotes the choice of agent i at time t and $n_t^A = \sum_i 1_{[a_t^i=A]}$ be the number of agents choosing action A at time t . Throughout the paper, we assume that each agent has to pick an action at every time t , i.e., the number of agents choosing action B is $n_t^B = N - n_t^A$.

2.1 Costs

We also assume that each action has a cost. For example, this could be total trip time or fuel consume. The total cost of action A at time t is a function of the number n_t^A of agents that pick A at time t . Let $c_A : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ denote the so-called cost function for action A . If n_t^A agents choose action A at time t , the cost to each of these agents is $c_A(n_t^A)$. We define c_B similarly to c_A . Figure 1a gives an example of these cost functions. The *social cost* $C(n_t^A)$ scales the costs of the two actions at time t with the proportions of agents taking the two actions, i.e.,

$$C(n_t^A) \triangleq \frac{n_t^A}{N} \cdot c_A(n_t^A) + \frac{n_t^B}{N} \cdot c_B(n_t^B). \quad (1)$$

¹Our results can be generalised to the case of arbitrarily many choices. The main change consists of replacing binomial probability distributions by multi-nomial ones.

The social cost corresponding to the cost functions of Figure 1a is shown in Figure 1b. We also define the *time-averaged social cost* as follows:

$$\hat{C}_T \triangleq \frac{1}{T} \sum_{t=1}^T C(n_t^A), \quad (2)$$

which will be used to illustrate the evolution of the social cost in simulations. The *social optimum* n^* of the social cost C is defined in the usual way:

$$n^* = \min_{n \in \{0,1,\dots,N\}} \frac{n}{N} \cdot c_A(n) + \frac{N-n}{N} \cdot c_B(N-n). \quad (3)$$

Given this basic setting, we wish to control these various costs by influencing individual agents' decisions. This is done using simple signalling schemes. We consider two specific schemes.

- The first one de-synchronises greedy agents with different actuation delays, i.e., delays between receiving a signal and taking a corresponding action.
- The second one de-synchronises agents with various levels of risk aversion.

We will later show that under reasonable assumptions on the behaviour of agents, these signalling schemes lead to desirable outcomes.

3 Related Work

Before proceeding, we now briefly mention some related work. Note that the related work spans many fields and an exhaustive survey is neither possible nor intended here.

Congestion problems are important in networking, control, operations research, and game theory [20, 9, 23, 12]. Our work proposes signalling solutions that reduce congestion by introducing a non-trivial amount of uncertainty through randomisation. Although it may seem surprising given the simplicity of the model, our results are unknown in the literature to the best of our knowledge.

In operations research and mathematical optimisation, congestion problems are often considered to be deterministic nonlinear resource allocation problems (*e.g.*, [20, 9, 23, 12]). In our work, however, the decision-maker only chooses the signals to send, instead of allocating the resources, *per se*.

Game theoretical models assume agents that act strategically, i.e., each agent has a non-negligible effect on the outcome. The study of one-shot congestion games has a long history, starting with [13]. The socially optimal and Nash equilibrium outcomes of congestion games have been compared in [14, 1], and [2]. These works show that when agents have full information, a resulting equilibrium outcome can incur much higher total congestion than a socially optimal outcome. An important distinction of our work is that, instead of a static game theoretical model, we consider a dynamic model evolving over time. The agents

do not observe other agents’ actions, nor do they act strategically: the only information they have about the congestion of the resources is obtained from signals broadcast by a central agent. We model a heterogeneous population of agents, i.e., a large number of agents with a variety of *fixed* policies. These agents do not act strategically in the fashion of price-taking participants in large markets, which reflects situations where each agent has very limited impact on the outcome as a whole. We take the view of the central agent, who has perfect knowledge of the congestion across resources up to the previous time step, but imperfect knowledge of the stochastic composition of the population of users. Our goal is to minimise the social welfare or total congestion: we study how the amount of information withheld can contribute to better social outcomes. In contrast to evolutionary game theory, where agents of different types interact and the population profile (or distribution) evolves over time, the population profile remains fixed in our setting, and only the action profile evolve over time.

Our random signals are reminiscent of perturbation schemes in repeated games, such as Follow-the-Perturbed-Leader [6], trembling-hand equilibrium [19], stochastic fictitious play [7], or the power of two choices [11]. When the agents use such randomised algorithms in their decision-making, the resulting demand process has been shown to behave rather well in theory (e.g., [11]), as well as in a number of applications, e.g., in parking [17], bike sharing [15], and charging electric vehicles [16]. However, we propose a signalling and guidance scheme that combines randomization and intervals.

In the transportation literature, [21] introduces the notion of equilibrium as the limit of the congestion distribution if it exists. [8] considers a number of notions of noisy signals and studies greedy policies and equilibria. Our approach is also related to signalling of parking space availability [18].

Our interval signalling scheme is reminiscent of the equilibrium outcome of [4] in the context of signalling games in economics (cf. [22] for an up-to-date survey). However, we consider the problem of optimal signalling in a dynamic system, whereas [4] considers Nash equilibrium signalling in single-shot games. For signalling games, [5] shows that more information does not generally improve the equilibrium welfare of agents. Their notion of “information,” due to [3], is however very different from ours.

Finally, let us note that there is also a related draft [10] by the present authors, which explores the notion of r -extreme interval signalling and optimisation over truthful interval signals.

4 Scalar Signalling

In this section, we model agents with different actuation delays – *i.e.*, delays between receiving a signal and taking a corresponding action – and introduce a signalling scheme with the aim of desynchronising these agents.

Notationwise, let H_t denote the congestion costs at time t :

$$H_t \triangleq (c_A(n_t^A), c_B(n_t^B)) \in \mathbb{R}^2.$$




Figure 2: Block diagram for scalar signalling.

We write $s_t^i \triangleq s_t^i(H_{t-1})$ denotes the signal that the central agent sends to agent i at every time t . For a fixed integer d , a *signalling scheme* is a sequence of mappings

$$\{s_t^i : \mathbb{R}^2 \rightarrow \mathbb{R}^d \mid i = 1, \dots, N, t = 1, 2, \dots\}. \quad (4)$$

4.1 σ -Scalar Signalling

We propose the following scalar signalling scheme, which corresponds to the case $d = 2$ in (4), with a parameter $\sigma > 0$. At every time step t , the central agent sends to each agent i a distinct signal $s_t^i := (y_t^{A,i}, y_t^{B,i}) \in \mathbb{R}^2$:

$$y_t^{A,i} \triangleq c_A(n_{t-1}^A) + w_t^{i,A}, \quad (5)$$

$$y_t^{B,i} \triangleq c_B(n_{t-1}^B) + w_t^{i,B}, \quad (6)$$

where $\{(w_t^{i,A}, w_t^{i,B}) : i = 1, \dots, N; t = 1, 2, \dots\}$ are independent and identically distributed (i.i.d.) zero-mean normal random variables such that $\mathbb{E}(w_t^{i,A} - w_t^{i,B})^2 = \sigma^2$ for all i and t .

Remark 1 (Random Signals). The signals $\{s_t^i\}$ in this section and the next are generated by the random variables $\{w_t^{i,A}, w_t^{i,B}\}$ in this section and $\{\nu_t, \eta_t\}$ in the next. In the rest of the paper, all expectations and probabilities are with respect to the distribution of these random signals.

In contrast to signalling schemes that solve an allocation problem to minimise the social cost and coerce each agent towards its assigned action, σ -scalar signalling is *scalar-truthful* in the following sense:

$$\mathbb{E}y_t^{A,i} = c_A(n_{t-1}^A), \quad \text{and} \quad \mathbb{E}y_t^{B,i} = c_A(n_{t-1}^B),$$

for all i and t , where the expectation is over the random variables $\{w_t^{i,X}\}$. We will show in Section 4.3 that it is possible to reduce the social cost by setting the parameter σ to a non-trivial value.

4.2 Agent Population and Policies

In response to the history of signals received prior to time t , every agent i takes action a_t^i . We assume that every agent acts based only on the signals, without considering the response of other agents to its own action. This is a reasonable assumption for three reasons. First, it is hard for the agent to obtain more information than the signal sent by the central agent. Second, the agents know that the signals received are truthful. Finally, the environment may be rapidly changing, making independent decisions difficult.

Formally, let S_t^i denote the history of signals received by agent i up to time t : $S_t^i \triangleq \{s_1^i, \dots, s_{t-1}^i\}$. Let \mathcal{S}_t denote the set of possible realisations of signal histories up to time t . A sequence of mappings from signal history to action, $\mathcal{S}_t \rightarrow \{A, B\}$ for $t = 1, 2, \dots$, is called a policy.

We let Ω denote the set of all possible types of agents. We introduce a probability measure μ over the set Ω , which describes the distribution of the population of agents into types. We assume that each type $\omega \in \Omega$ is associated with a policy, and every agent of type ω follows the policy $\pi^\omega : \mathcal{S}_t \rightarrow \{A, B\}$. For a population of N agents, we assume that μ is such that $\mu(O)N$ is a non-negative integer for all $O \subseteq \Omega$. Consequently, $\mu(\omega)N$ denotes the number of agents with policy π^ω . The following assumption simplifies the analysis.

Assumption 1. *The true distribution μ is known to the central agent.*

It is also reasonable that the central agent is capable of estimating the distribution of agents μ using statistical estimation techniques. Estimating distribution of agents is possible by Bayesian or maximum a-posteriori methods. If the estimation is not accurate, the central agent can always fall back to sending no signal, which creates no more congestion than there already is.

4.2.1 π^k -policies

In the case of scalar signalling, we assume that the set of possible types is $\Omega = \{1, 2, \dots\}$. Each agent of type $k \in \Omega$ has an actuation delay of k time steps, i.e., it acts at time t upon the signal s_{t-k}^i received at the time step $t - k$ – if the latter is defined. This effectively models the delay between when an agent makes the decision which action to take, and when it contributes to congestion. These delays do not come from communication, but from the agent-specific

delay between making a decision and causing congestion. For example, a driver may decide to take route A at time $t - k$, but begin the journey only at time t , resulting in the delay between deciding on a route and reaching a particular congested road segment. Another example is the computational delay between receiving stock market information and deciding to invest in a stock.

More precisely, in response to the scalar signalling scheme with parameter σ , i.e., $s_t^i := (y_t^{A,i}, y_t^{B,i})$, each agent i of type k acts according to the following policy π^k :

$$\begin{aligned} a_t^i &= \pi^k(S_t^i) = \pi^k(s_{t-k}^i) \\ &= \begin{cases} \arg \min_{X=A,B} y_{t-k}^{X,i}, & \text{if } t \geq k + 1, \\ A, & \text{otherwise,} \end{cases} \end{aligned} \quad (7)$$

i.e., the agent chooses greedily the action with the smallest cost signal at time $t - k$.

4.3 Guarantees

In this section, we characterise the social cost resulting from σ -scalar signalling and a population of agents following policies $\{\pi^k : k = 1, 2, \dots\}$. Furthermore, we show that the expected distribution of agents between the actions converges for every initial condition.

For clarity of exposition, we first consider a homogeneous population. We analyse the expected value and concentration property of the social cost, and show the convergence of each agent's action profile. Our analysis methods can be used to obtain corresponding guarantees for the heterogeneous population cases. We also illustrate our main theoretical guarantees with simulations.

4.3.1 Homogeneous Population

In this section, we consider a homogeneous population of agents with the policy π^1 , i.e., the case $\mu(1) = 1$. In other words, we assume that every agent i acts according to π_1 .

First, we focus on the next-step outcome and establish closed-form expressions for the expectation of the number of agents taking action A .

Lemma 1 (Conditional Distribution of n_t^A). *Suppose that n_{t-1}^A takes a fixed value $n \in \{1, \dots, N\}$ and that all agents are of policy π^1 . Consider an arbitrary time step $t > 1$ and arbitrary n_{t-1}^A . Let $p_{\sigma,n} \triangleq \Phi_{\sigma}(c_B(N - n) - c_A(n))$. Then, we have*

$$\mathbb{P}(n_t^A = m \mid n_{t-1}^A = n) = \binom{N}{m} p_{\sigma,n}^m (1 - p_{\sigma,n})^{N-m}.$$

Proof. Recall that $y_1^{A,i} := c_A(n_t^A) + w_i$, and $y_1^{B,i} := c_B(N - n_t^A) + w_i$. Notice that a_t^i is a random variable. The number of agents taking action A at time t

is:

$$n_t^A = \sum_{i=1}^N 1_{[a_t^i=A]}.$$

Clearly, for a fixed t , $\{w_{i,t}^A - w_{i,t}^B : i = 1, \dots, N\}$ is i.i.d. with normal distribution $N(0, \sigma^2)$. Hence, for a fixed t , $\{1_{[a_t^i=A]} : i = 1, \dots, N\}$ are i.i.d. Bernoulli random variables, each with parameter $\mathbb{P}(a_t^i = A)$.

Let Φ_σ denote the tail probability function of the distribution $N(0, \sigma^2)$. Observe that

$$\begin{aligned} \mathbb{P}(a_t^i = A) &= \mathbb{P}(y_t^{A,i} < y_t^{B,i}) \\ &= \mathbb{P}(c_A(n_{t-1}^A) + w_{i,t}^A < c_B(N - n_{t-1}^A) + w_{i,t}^B) \\ &= \mathbb{P}(w_{i,t}^A - w_{i,t}^B < c_B(N - n_{t-1}^A) - c_A(n_{t-1}^A)) \\ &= \Phi_\sigma(c_B(N - n) - c_A(n)) = p_{\sigma,n}, \end{aligned}$$

where we used the assumption that $n_{t-1}^A = n$.

Since n_t^A is a sum of Bernoulli random variables, the claim follows by the binomial probability mass function. \square

Next, we can derive a closed-form expression for the expectation of the next-step social cost.

Theorem 1 (Expected Next-step Social Cost). *Let us take the assumptions as in Lemma 1. For every parameter $\sigma > 0$, and allocation $n \in \{1, \dots, N\}$ to action A , we have*

$$\frac{\mathbb{E}C(n_t^A)}{C(n)} = \sum_{m=-\infty}^{\infty} \frac{C(m)}{C(n)} \binom{N}{m} p_{\sigma,n}^m (1 - p_{\sigma,n})^{N-m},$$

where the expectation is conditioned on $n_{t-1}^A = n$.


Proof. Observe that

$$\begin{aligned} \mathbb{E}C(n_t^A) &= \sum_{m=-\infty}^{\infty} C(m) \mathbb{P}(n_t^A = m) \\ &= \sum_{m=-\infty}^{\infty} C(m) \binom{N}{m} p_{\sigma,n}^m (1 - p_{\sigma,n})^{N-m}. \end{aligned} \tag{8}$$

The claim follows. \square

Recall that n^* denotes the socially optimal allocation of agents to action A . Figure 3a² illustrates the expected social cost of Theorem 1 as a function of the

²In Figure 3a and throughout the paper, all error-bar plots are averaged over 100 simulations with error bars corresponding to one standard deviation.



(a) Next-step social cost $C(n_t^A)$ versus parameter σ with $n_{t-1}^A = n^*$. Observe that the social cost is minimum at $\sigma = 0.3$ and that the variance decreases as we approach the minimum.

(b) Time-averaged social cost \hat{C}_{30} versus parameter σ with $n_1^A = n^*$. The dashed horizontal line is at the unique optimum.

Figure 3: σ -scalar signalling, homogeneous population of $N = 40$, and the cost functions of Figure 1a.

parameter σ , along with the simulated mean and standard deviation of $C(n_t^A)$. In particular, observe that by setting $\sigma = 0.3$, the next-step social cost $C(n_t^A)$ achieves the theoretical minimum social cost $C(n^*)$ of Figure 3a, provided that the population distribution is optimal at the current time step ($n_{t-1}^A = n^*$). Furthermore, Figure 3b shows the time-averaged social cost \hat{C}_T can achieve values close to the theoretical minimum social cost by setting the parameter σ appropriately. Hence, σ -scalar signalling is useful even when we do not satisfy the condition $n_{t-1}^A = n^*$ for many time indices t .

Remark 2 (Optimal σ). It is clear from Figure 3a and by substituting $p_{\sigma,n} \triangleq \Phi_\sigma(c_B(N-n) - c_A(n))$ into Theorem 1 that we can readily minimise the next-step expected social cost by optimising the amount of noise σ .

Next, we quantify the concentration of the social cost around its expected value, which is already hinted at in the error bars of Figure 3a.

Theorem 2 (Concentration of Next-step Social Cost). *Suppose that the assumptions of Theorem 1 and that the functions c_A and c_B are Lipschitz with constant L . We have*

$$\mathbb{P}\left(\left|\frac{C(n_t^A)}{C(n^*)} - \frac{\mathbb{E}C(n_t^A)}{C(n^*)}\right| \geq \varepsilon\right) \leq 2 \exp\left(-\frac{2\varepsilon^2 C(n^*)}{2(L+1)}\right),$$

where the probability is conditioned on $n_{t-1}^A = n^*$.

Proof. Let

$$X_i = 1_{[a_t^i = A]}, \quad \text{for all } i = 1, \dots, N,$$

so that $n_t^A = \sum_i X_i$. Let f be such that

$$\begin{aligned} f(X_1, \dots, X_N) &= C(n_t^A)/C(n^*) \\ &= \frac{1}{C(n^*)} \left[\frac{\sum_i X_i}{N} \cdot c_A \left(\frac{\sum_i X_i}{N} \right) \right. \\ &\quad \left. + \frac{N - \sum_i X_i}{N} \cdot c_B \left(\frac{N - \sum_i X_i}{N} \right) \right]. \end{aligned}$$

Let $x_1, \dots, x_N, \hat{x}_i \in \{0, 1\}$. Observe that since c_A and c_B take values in $[0, 1]$ and are L -Lipschitz by assumption, by simple algebra, we obtain

$$\begin{aligned} &\sup |f(x_1, \dots, x_N) - f(x_1, \dots, \hat{x}_i, \dots, x_N)| \\ &\leq \frac{2(1/N + L/N)}{C(n^*)}, \end{aligned}$$

for all i . Hence, by McDiarmid's inequality, we obtain the claim. \square

Next, we consider the normalised random process n_t^A/N for $t = 0, 1, 2, \dots$. In particular, we consider the expectation of that process, i.e., the sequence

$$x_t^i \triangleq \mathbb{P}(a_t^i = A) = \mathbb{E} \frac{n_t^A}{N}, \quad t = 1, 2, \dots$$

We show that this sequence converges. In order words, the action profile of every agent converges under σ -scalar signalling if every agent follows the π^1 policy (7).

Theorem 3 (Convergence). *Suppose that the assumptions of Theorem 1 hold. For every initial condition $x_1^i = \mathbb{P}(a_1^i = A)$, the sequence $x_t^i = \mathbb{P}(a_t^i = A)$ converges as $t \rightarrow \infty$.*

Proof. Observe that, for all $t > 2$, we have

$$x_t^i = \mathbb{P}(a_t^i = A) \tag{9}$$

$$= \mathbb{P}(w_{i,t}^A - w_{i,t}^B < c_B(N - n_{t-1}^A) - c_A(n_{t-1}^A)) \tag{10}$$

$$= \sum_m \mathbb{P}(w_{i,t}^A - w_{i,t}^B < c_B(N - m) - c_A(m)) \mathbb{P}(n_{t-1}^A = m) \tag{11}$$

$$= \sum_m \Phi_\sigma(c_B(N - m) - c_A(m)) \mathbb{P}(n_{t-1}^A = m), \tag{12}$$

where

$$\mathbb{P}(n_{t-1}^A = m) = \binom{N}{m} \mathbb{P}(a_{t-1}^i)^m (1 - \mathbb{P}(a_{t-1}^i))^{N-m} \tag{13}$$

$$= \binom{N}{m} (x_{t-1}^i)^m (1 - x_{t-1}^i)^{N-m}. \tag{14}$$




Figure 4: σ -scalar signalling, homogeneous population of $N = 40$, and the cost functions of Figure 1a.

By combining (12) and (13), we obtain a function $f : \mathbb{R} \rightarrow \mathbb{R}$ that describes the iterated function process $\{x_t^i : t = 1, 2, \dots\}$:

$$x_t^i = f(x_{t-1}^i), \quad \text{for all } t.$$

It is easy to verify that there exists an $\ell < 1$ such that for all x, y ,

$$|f(x) - f(y)| \leq \ell |x - y|.$$

Hence, f is a contraction and the claim follows from Banach fixed-point theorem. \square

To illustrate Theorem 3, Figure 4a shows the process x_t^i for various values of σ , whereas Figure 4b shows the related random process $\{n_t^A/N\}$. Observe that for the parameter value $\sigma = 0.6$, the process $\{n_t^A/N\}$ converges to values close to the optimal population profile $n^*/N = 0.2$ (cf. Figure 1b), whereas for other parameter values, we see flapping in the population profiles and values away from the optimum.

As a corollary of Theorem 3, we obtain the following.

Corollary 1. *The random processes n_t^A/N and $C(n_t^A)$ converge in distribution as $t \rightarrow \infty$.*

Remark 3 (Limit of $C(n_t^A)$). It follows from Theorem 3 that the limit distribution of the random process $C(n_t^A)$ depends on the parameter σ , which can be optimised.

4.3.2 Heterogeneous Population

In this section, we consider a distribution $\mu = (\mu_1, \dots, \mu_K)$ of agents over the policies π^1, \dots, π^K , i.e., with $\mu_j N$ agents with delays j for $j = 1, \dots, K$. We first derive the following analogue of Lemma 1.

Lemma 2 (Conditional Distribution of n_t^A). *Suppose that Assumption 1 holds. Consider an arbitrary time step $t > K$ and suppose that $n_{t-k}^A = n_0$ for all*




Figure 5: σ -scalar signalling on the cost functions of Figure 1a with $N = 40$ agents.

$k = 1, \dots, K$. Let $p_\sigma \triangleq \Phi_\sigma(c_B(N - n_0) - c_A(n_0))$. Then, we have

$$\begin{aligned} \mathbb{P}(n_t^A = m \mid n_{t-k}^A = n_0 \text{ for } k = 1, \dots, K) \\ = \sum_{|\alpha|=m} \binom{m}{\alpha} \prod_{j=1}^K \binom{\mu_j N}{\alpha_j} p_\sigma^{\alpha_j} (1 - p_\sigma)^{\mu_j N - \alpha_j}, \end{aligned}$$

where $\alpha = (\alpha_1, \dots, \alpha_K)$ denotes a multi-index.

Proof. Let $n_{j,t}^A$ denote the number of agent of type j who choose action A at time t . By the same argument as Lemma 1, we have

$$\mathbb{P}(n_{j,t}^A = m) = \binom{\mu_j N}{m} p_\sigma^m (1 - p_\sigma)^{\mu_j N - m}, \quad j = 1, \dots, K.$$

By definition, we have

$$\begin{aligned} \mathbb{P}(n_t^A = m) &= \mathbb{P}\left(\sum_{j=1}^K n_{j,t}^A = m\right) \\ &= \sum_{|\alpha|=m} \binom{m}{\alpha} \mathbb{P}(n_{1,t}^A = \alpha_1) \dots \mathbb{P}(n_{K,t}^A = \alpha_K). \end{aligned}$$

The claim follows by algebra. \square

Using Lemma 2, we can readily derive analogues of the expected next-step social cost (Theorem 1), the concentration of social cost (Theorem 2), and the convergence of action profiles (Theorem 3) guarantees for heterogeneous populations. Figures 5a illustrates the convergence of the population profile n_t^A/N with a heterogeneous population under σ -scalar signalling. Observe that, as in Figure 4b, a cyclic flapping behaviour is visible at low values of δ , e.g., $\delta = 0.1$.

An empirical study of the time-averaged social cost \hat{C}_t shows a dependence on the parameter σ similar to the next-step social cost $C(n_t^A)$, which can be seen by comparing Figure 5b to Figure 3a.




Figure 6: Block diagram for interval signalling.

5 Interval Signalling

In this section, we model agents with various levels of risk aversion, and present a broadcast signalling scheme that desynchronises them. Each signal is composed of pairs of endpoints, hence, the name interval signalling. The overall closed-loop system composed of the central agent and the agents is shown in Figure 6.

5.1 (δ, γ) -Interval signalling

In this section, we present a signalling scheme that broadcast the same signal to all agents. Let δ and γ denote non-negative constants. In the (δ, γ) -interval signalling scheme, a central agent broadcasts to all agents the same signal $s_t^i = s_t$ for all i . This is in contrast to σ -scalar signalling, where the central agent sends a distinct signal to every agent i . For every time step t , the signal $s_t = (\underline{u}_t^A, \bar{u}_t^A, \underline{u}_t^B, \bar{u}_t^B) \in \mathbb{R}^4$ is defined as follows³:

$$\underline{u}_t^A \triangleq c_A(n_{t-1}^A) + \nu_t - \delta/2, \quad \bar{u}_t^A \triangleq c_A(n_{t-1}^A) + \nu_t + \delta/2, \quad (15)$$

$$\underline{u}_t^B \triangleq c_B(n_{t-1}^B) + \eta_t - \gamma/2, \quad \bar{u}_t^B \triangleq c_B(n_{t-1}^B) + \eta_t + \gamma/2. \quad (16)$$

where ν_t and η_t are i.i.d. uniform random variables with supports:

$$\text{Supp}(\nu_t) = [-\delta/2, \delta/2], \quad \text{Supp}(\eta_t) = [-\gamma/2, \gamma/2].$$

Figure 7 illustrates the interval $[\underline{u}_t^A, \bar{u}_t^A]$.

³This corresponds to the case $d = 4$ in (4).




Figure 7: An illustration of interval signalling.

By construction, it is clear that (δ, γ) -interval signalling scheme has the following property, which we call *interval-truthfulness*:

$$c_A(n_{t-1}^A) \in [\underline{u}_t^{A,i}, \bar{u}_t^{A,i}] \text{ and } c_B(n_{t-1}^B) \in [\underline{u}_t^{B,i}, \bar{u}_t^{B,i}]$$

with probability 1 for all i and t . We will show in Section 5.2 that it is also possible to pick the parameters (δ, γ) to reduce the social cost.

5.1.1 π^λ -policies

In the case of interval signalling, we consider a set of agent types $\Omega = [0, 1]$, where the type captures a level of risk aversion. Recall that every agent receives the same interval signal $s_t := (\underline{u}_t, \bar{u}_t, \underline{v}_t, \bar{v}_t)$. In response, we assume that each agent of type λ follows the policy π^λ :

$$a_t^i = \pi^\lambda(S_t^i) = \pi^\lambda(s_t) := \arg \min_{X \in \{A, B\}} \lambda \underline{u}_t^X + (1 - \lambda) \bar{u}_t^X. \quad (17)$$

These policies naturally model a notion of risk aversion. Observe that for $\lambda = 0$, policy π^0 models a risk-averse agent, who makes decisions based solely on upper endpoints \bar{u}_t^A and \bar{u}_t^B of the respective intervals. Similarly, π^1 and $\pi^{1/2}$ model risk-seeking and risk-neutral agents. Although the extremes may be rare, it seems plausible that people combine the optimistic and pessimistic views in this fashion. The following example illustrates these policies.

Example 1. Suppose that at time $t = 1$, we have $c_A(n_1^A) = 1$ and $c_B(n_1^B) = 1$. For $\delta = 1, \gamma = 0.5$, the interval signal $s_2 := (\underline{u}_2, \bar{u}_2, \underline{v}_2, \bar{v}_2)$ at time 2, could take the realisation of $(0.5, 1.5, 0.8, 1.3)$. The risk-seeking agent with $\lambda = 1$ would pick A , whereas the risk-averse agent with $\lambda = 0$ would pick B .

5.2 Guarantees

In this section, we consider (δ, γ) -interval signalling and π^λ -policies. For signalling scheme that broadcasts the same signal to all agents, the outcome corresponding to homogeneous populations is trivial since all agents would pick the same action. Hence, we consider directly a heterogeneous population over a set of types $[0, 1]$ and a population distribution μ .

First, we derive a lemma for the conditional probability of taking action A . The main result of this section is an expression to evaluate the expected next-step social cost for a μ -heterogeneous population under (δ, γ) -interval signalling.

Lemma 3 (Conditional Probability of Action A). *Let $F_{\delta,\gamma,n}(\lambda) \triangleq \mathbb{P}(\pi^\lambda(s_t) = A \mid n_{t-1}^A = n)$ denote, for every $\lambda \in \Omega$, the probability of an agent with policy π^λ choosing action A conditioned on the realisation of the random variable n_{t-1}^A . We have*

$$F_{\delta,\gamma,n}(\lambda) = \mathbb{P}(\nu_t + \eta_t > c_A(n) + \delta(1/2 - \lambda) - c_B(N - n) - \gamma(1/2 - \lambda)),$$

where ν_t and η_t are independent uniform random variables with supports $[-\delta/2, \delta/2]$ and $[-\gamma/2, \gamma/2]$.

Proof. For every $\lambda \in \Omega$, the probability of an agent with policy π^λ choosing action A is

$$\begin{aligned} \mathbb{P}(\pi^\lambda(s_t) = A) &= \mathbb{P}(\lambda \underline{u}_t^A + (1 - \lambda) \bar{u}_t^A < \lambda \underline{u}_t^B + (1 - \lambda) \bar{u}_t^B) \\ &= \mathbb{P}(c_A(n_{t-1}^A) + \delta(1/2 - \lambda) - c_B(N - n_{t-1}^A) - \gamma(1/2 - \lambda) < -\nu_t + \eta_t), \end{aligned}$$

where the first equality follows by definition of the policy π^λ , and the second equality follows by the definition of (δ, γ) -interval signalling and simple algebra. The claim follows from the fact that $-\nu_t$ has the same distribution as ν_t by definition. \square

Theorem 4 (Expected Next-Step Social Cost). *Suppose that the assumptions of Lemma 3 hold with $n = n^*$. Let $p_{\delta,\gamma,n^*} \triangleq \int_{\lambda \in [0,1]} F_{\delta,\gamma,n^*}(\lambda) \mu(d\lambda)$. We have*

$$\frac{\mathbb{E}C(n_t^A)}{C(n^*)} = \sum_{m=-\infty}^{\infty} \frac{C(m)}{C(n^*)} \binom{N}{m} p_{\delta,\gamma,n^*}^m (1 - p_{\delta,\gamma,n^*})^{N-m}.$$


Proof. By the definition $F_{\delta,\gamma,n^*}(\lambda) \triangleq \mathbb{P}(\pi^\lambda(s_t) = A \mid n_{t-1}^A = n^*)$, we have

$$\mathbb{P}(n_t^A = m \mid n_{t-1}^A = n^*) = \binom{N}{m} p_{\delta,\gamma,n^*}^m (1 - p_{\delta,\gamma,n^*})^{N-m}.$$

The claim follows by the same argument as the proof of Theorem 1 – cf. (8). \square


Figure 8 presents in heat-map form the dependence of the next-step social cost $C(n_{t+1}^A)$ on the parameters δ and γ of the interval signalling scheme. Clearly, the optimal (δ, γ) -values are a non-trivial region bounded away from 0 and ∞ . The simulated next-step social cost $C(n_{t+1}^A)$ on Figure 8a coincides with the expected value of $C(n_{t+1}^A)$ computed using Theorem 4 on Figure 8b.

Remark 4 (Concentration of $C(n_t^A)/C(n^*)$). In the case of (δ, γ) -interval signalling, the social cost ratio $C(n_t^A)/C(n^*)$ is concentrated around the mean in a similar fashion to Theorem 2.



(a) Simulated next-step social cost $C(n_{t+1}^A)$ versus δ and γ . (b) Expected $C(n_{t+1}^A)$ computed using Theorem 4 as a function of δ and γ

Figure 8: (δ, γ) -interval signalling, cost functions of Figure 1a, and heterogeneous population with $N = 40$ and μ uniform over $\{1/N, \dots, 1\}$.



(a) The processes $\{n_t^A/N\}$ over time for $N = 40$ agents, $\delta = 0.1$, and $\gamma = 0.2$. (b) Time-averaged social cost \hat{C}_{30} versus δ for a fixed $\gamma = 0.2$, uniform distribution μ over $[0, 1]$, and $N = 40$. The dashed horizontal line is at the optimal δ .

Figure 9: (δ, γ) -interval signalling on the cost functions of Figure 1a

5.3 Long-Run Behaviour

In this section, we study empirically the long-run behaviour of the state process n_t^A/N under (δ, γ) -interval signalling. Figure 9a shows the simulated the time evolution of n_t^A/N . As in Figure 4b, a severe flapping effect is exhibited at a low value of δ , e.g., $\delta = 0.5$. For higher values of δ , there are hints of convergence.

Figure 9b shows the dependence of the time-averaged social cost \hat{C}_T on one parameter δ , with the other parameter fixed at $\gamma = 0.2$. Observe that the dependence is similar to that exhibited in the case of σ -scalar signalling (cf. 3a); moreover, at the optimal value of the parameter δ , the social cost – although random – is close to the theoretical minimum (cf. 1b).

6 Conclusion

Our analysis and simulations give quantitative guidance on how to design signalling schemes in order to reduce the social cost. In particular, if we can estimate the true population mixture, e.g., using statistical estimation techniques, then we can optimise the signalling scheme to minimise the social cost. Although our analysis focuses on a congestion problem with two resources, the results can be extended to an arbitrary number of resources, e.g., by replacing the binomial distributions by a multinomial distribution in the analysis.

In contrast to signals that report past outcomes as deterministic numbers, which correspond to singular probability distributions, we propose random signals with non-singular distributions. When agents mistake deterministic signals for precise predictions of the future, their actions lead to the cyclic flapping behaviour. We show that agents who receive random signals with non-singular distributions exhibit cyclic behaviour to a lesser degree, which can improve the social cost.

Regarding future work, we pose several new questions that can be studied by extending our model. What happens when population size and composition change over time? This can be modeled by a Markovian sequence of random variables corresponding to agent types. How do the optimal parameter values for σ, δ, γ depend on those changes? This can be done by employing state-space control techniques to obtain these optimal values. Would our results hold if every agent had a distinct cost function, so as to model the priority given to some agents, e.g., when evacuating the young and elderly? This requires a new notion of consistency or truthfulness that is specialised to individual agents, as well as a new analysis.

Acknowledgement

This work was supported in part by the EU FP7 project INSIGHT under grant 318225, and in part by Science Foundation Ireland grant 11/PI/1177.

References

- [1] E. Anshelevich, A. Dasgupta, J. Kleinberg, and E. Tardos. The price of stability for network design with fair cost allocation. In *Foundations of Computer Science*, 2004.
- [2] M.-F. Balcan, A. Blum, and Y. Mansour. The price of uncertainty. In *Electronic Commerce*, 2009.
- [3] D. Blackwell. Comparison of experiments. In J. Neymann, editor, *Proceedings of the Second Berkeley Symposium on Mathematical Statistics and Probability*, pages 93–102, 1951.
- [4] V. P. Crawford and J. Sobel. Strategic information transmission. *Econometrica*, 50(6):1431–1451, 1982.
- [5] J. R. Green and N. L. Stokey. A two-person game of information transmission. *Journal of Economic Theory*, 135(1):90 – 104, 2007.
- [6] J. Hannan. Approximation to Bayes risk in repeated play. In *Contributions to the Theory of Games*, volume 3, pages 97–139. Princeton University Press, 1957.

- [7] J. C. Harsanyi. Games with randomly disturbed payoffs: A new rationale for mixed-strategy equilibrium points. *International Journal of Game Theory*, 2(1):1–23, 1973.
- [8] J. L. Horowitz. The stability of stochastic equilibrium in a two-link transportation network. *Transportation Research Part B: Methodological*, 18(1):13–28, 1984.
- [9] T. Ibaraki and N. Katoh. *Resource Allocation Problems: Algorithmic Approaches*. MIT Press series in the foundations of computing. Mit Press, 1988.
- [10] J. Marecek, R. Shorten, and J. Y. Yu. Congestion management by interval signaling. *CoRR*, abs/1404.2458, 2014.
- [11] M. Mitzenmacher. The power of two choices in randomized load balancing. *Parallel and Distributed Systems, IEEE Transactions on*, 12(10):1094–1104, Oct 2001.
- [12] M. Patriksson. A survey on the continuous nonlinear resource allocation problem. *European Journal of Operational Research*, 185(1):1–46, 2008.
- [13] R. W. Rosenthal. A class of games possessing pure-strategy nash equilibria. *International Journal of Game Theory*, 2(1):65–67, 1973.
- [14] T. Roughgarden and E. Tardos. How bad is selfish routing? *Journal of the ACM*, 49(2):236–259, 2002.
- [15] A. Schlote, B. Chen, and R. Shorten. On closed-loop bicycle availability prediction. *Intelligent Transportation Systems, IEEE Transactions on*, PP(99):1–7, 2014.
- [16] A. Schlote, F. Hausler, T. Hecker, A. Bergmann, E. Crisostomi, I. Radusch, and R. Shorten. Cooperative regulation and trading of emissions using plug-in hybrid vehicles. *Intelligent Transportation Systems, IEEE Transactions on*, 14(4):1572–1585, Dec 2013.
- [17] A. Schlote, C. King, E. Crisostomi, and R. Shorten. Delay-tolerant stochastic algorithms for parking space assignment. *Intelligent Transportation Systems, IEEE Transactions on*, 15(5):1922–1935, Oct 2014.
- [18] A. Schlote, C. King, E. Crisostomi, and R. Shorten. Delay-tolerant stochastic algorithms for parking space assignment. *Intelligent Transportation Systems, IEEE Transactions on*, 15(5):1922–1935, Oct 2014.
- [19] R. Selten. Reexamination of the perfectness concept for equilibrium points in extensive games. *International Journal of Game Theory*, 4(1):25–55, 1975.
- [20] Y. Sheffi. *Urban transportation networks: equilibrium analysis with mathematical programming methods*. 1985.

- [21] M. Smith. The existence, uniqueness and stability of traffic equilibria. *Transportation Research Part B: Methodological*, 13(4):295–304, 1979.
- [22] J. Sobel. Giving and receiving advice. In D. Acemoglu, M. Arellano, and E. Dekel, editors, *Advances in Economics and Econometrics*, 2013.
- [23] S. M. Stefanov. *Separable Programming: Theory and Methods*. Applied Optimization. Springer, 2001.